CSO The Resource for Security Executives
csoonline.com Home | Magazine | Newsletters | Career | Online Features | Resources | Search
April 2003 CSO Magazine

machine shop

Technologies, Tools and Tactics
Edited by Derek Slater

Hard-Disk Risk

Are all those old hard drives you're getting rid of free of important company data? Don't be so sure. **BY SIMSON GARFINKEL**

A FEW YEARS AGO, when I was in Silicon Valley with nothing to do, I stopped by one of the valley's famed stores that sell used and "recycled" computers. In the store's front were used minicomputers, workstations, terminals and lots of old PCs that had all seen better days. Then I noticed that the store was selling used hard drives as well. A 10GB drive could be had for just $30—quite a bargain at the time.

"You clear the information off these drives before you sell them?" I asked innocently.

"Absolutely," said the man behind the counter. "I do it myself. We run FDisk on every drive. There's no way to get back the information after you do that."

Really? Turns out he was wrong. Running Windows FDisk on a 10GB drive overwrites only 0.01 percent of the drive's sectors. Although Windows doesn't give you any tools for recovering the data afterward, many such tools are currently on the market (for descriptions of those tools, see "Tools of Evidence," Machine Shop, March 2003).

But the real treasure trove that day wasn't on the store's display shelves; it was in the warehouse. The cavernous space out back had several shelves stacked high with old hard drives, each $5, "as is and untested," according to the sign. In other words, nobody had even run FDisk on those drives. Pop one into a computer, and you could recover the previous owner's files simply by running XCopy.

I bought 20 of them.

I took the drives home and started my own forensic analysis. Several of the drives had source code from high-tech companies. One drive had a confidential memorandum describing a biotech project; another had internal spreadsheets belonging to an international shipping company.

Since then, I have repeatedly indulged my habit for procuring and then analyzing secondhand hard drives. I bought recycled drives in Bellevue, Wash., that had internal Microsoft e-mail (somebody who was working from home, apparently). Drives that I found at an MIT swap meet had financial information on them from a

Boston-area investment firm. Last summer, I started buying drives en masse on eBay.

In all, I bought and analyzed the content of more than 150 drives with the help of Abhi Shelat, another graduate student at MIT's Laboratory for Computer Science. We found that between one-third and one-half of the drives still had significant amounts of confidential data, even though many had been through a Format or FDisk operation. On another third, someone had deleted the document files but left the applications behind. It was a simple matter to undelete the data files and retrieve their secrets as well.

In fact, only 10 percent of the drives I purchased had been properly sanitized.

Much of the data we found was truly shocking. One of the drives once lived in an ATM. It contained a year's worth of financial transactions—including account numbers and withdrawal amounts—from a organization that had a legal requirement to not divulge such information. Two other drives contained more than 5,000 credit card numbers—it looked as if one had been inside a cash register. Another had e-mail and personal financial records of a 45-year-old fellow in Georgia. The man is divorced, paying child support and dating a woman he met in Savannah. And, oh yeah, he's really into pornography.

Abhi and I published our findings earlier this year in *IEEE Security and Privacy* journal. The story got a lot of media attention. It seems that many people have heard that some used computers still have confidential information on their hard drives, but few suspected the scale of the problem.

Suds for Your Hard Drive

So what's to be done?

Perhaps the saddest observation in our story is that erasing information from hard drives is not difficult—with a little bit of Web searching, we found more than 50 programs that purport to clean your hard drive so that the information on it cannot be recovered using even the most advanced technical means. One program costs more than $1,000, but some cost only $20 or $30, while still others are free. All of the programs do more or less the same thing: They repeatedly overwrite the blocks on your computer's hard drive with random bit patterns, completely obscuring the information that was previously there.

These so-called disk sanitizers actually come in two varieties. The first is programs that promote themselves as file shredders, secure erasers or slack-space sanitizers, designed to be used on a running computer system. They overwrite blocks on your disk that aren't actively being used to store files but might have been used in the past for file storage. These programs, such as SecureClean from AccessData, assure that deleted files are no longer recoverable. The best will sanitize other kinds of telltale privacy leaks, including browser caches, temporary files and certain kinds of cookies.

The second kind of program will completely erase the contents of a disk—just the thing when you want to upgrade the PCs in the accounting department and redeploy

them on reception desks throughout your enterprise. The programs, properly called disk sanitizers but sometimes called disk shredders, repeatedly overwrite every block of a disk drive, then fill the drive with zeros.

The best disk sanitizers come on a bootable floppy or CD-ROM. You insert the removable media into the computer to be wiped clean, boot the computer and verify your intentions to the program. It does the rest. Clearly, these programs can be dangerous in the hands of a disgruntled employee—one reason it's always a good idea to restrict physical access to your most important systems. One disk sanitizer I'm particularly fond of is called Autoclave. You can download it from *staff.washington.edu/jdlarios/autoclave*, write it to a floppy and go to town.

But the study that Abhi and I did shows that many organizations are simply not taking the problem seriously.

One key reason for today's poor disk sanitization practices is that it's very difficult to tell the difference between a disk that has been properly sanitized and one that's simply been reformatted. Both look blank to the untrained technician—you need forensic tools to tell the difference. You also need to put the drive in a working computer. So simply checking to see if a disk is sanitized can be prohibitively expensive in many cases.

Another reason, we suspect, is that most people don't appreciate the risk—the used-computer market is literally awash with personal information from businesses and individuals, yet there are relatively few cases of that information being used for nefarious purposes.

The used-computer market is literally awash with personal information from businesses and individuals, yet there are relatively few cases of that information being used for nefarious purposes.

Is data left on salvaged hard drives a problem for the typical CSO? I think it is. We spend so much time and money trying to protect the information on our computers, it's utterly irresponsible for us to then just throw it out. Why should the confidentiality of data in your organization depend on the good intentions of a person who buys one of your used drives?

Search and Recovery

This whole world of disk sanitization can be very off-putting to the average CSO. Many people maintain that shadowy organizations such as the National Security Agency can retrieve data from a hard drive even after that data has been overwritten with a random pattern. Some say that you need to overwrite a hard drive not once, but seven or even 22 times.

Such lore has even made its way into the disk sanitization programs. SuperScrubber from Jiiva, one of the few Macintosh data sanitization products, offers five so-called security levels: Simple (not secure), Simple + Verify (not secure), Strong, Military and Paranoid. Why in heaven's name would a security professional use a security program in a manner that the program itself claims is not secure? Such attitudes and programs make the task of erasing hard drives seem so daunting that many people are apparently scared away. Why try to solve a problem that's basically unsolvable?

In fact, there is no unclassified evidence that data on a modern hard drive can be recovered after it has been overwritten with just a single pass of random information. Some have made such claims, but no such recovery has ever been demonstrated in public. Today's hard drives are specifically designed not to work that way. When you save a new version of a Microsoft Word file on your hard drive, for instance, you want to get the new—not the old—version.

A growing number of businesses offer to properly sanitize, refurbish and reload your computers with "clean" software before the machines are repurposed within your organization or sold. Although outsourcing sounds attractive, I'm concerned that it is exceptionally difficult to audit those companies and make sure they are actually deleting your data.

In the end, preventive technology is a better solution to the sanitization problem. If you use an encrypted file system, you can sanitize a disk simply by erasing the key. I'd like to see that sort of technology built in to hard drives. Or better, perhaps someday soon, all disk drives will come with a self-destruct feature—just like *Star Trek*'s *Enterprise* did!

Simson Garfinkel, CISSP, is a technology writer based in the Boston area. He is also CTO of Sandstorm Enterprises, an information warfare software company. He can be reached at *machineshop@cxo.com*.